Section 3

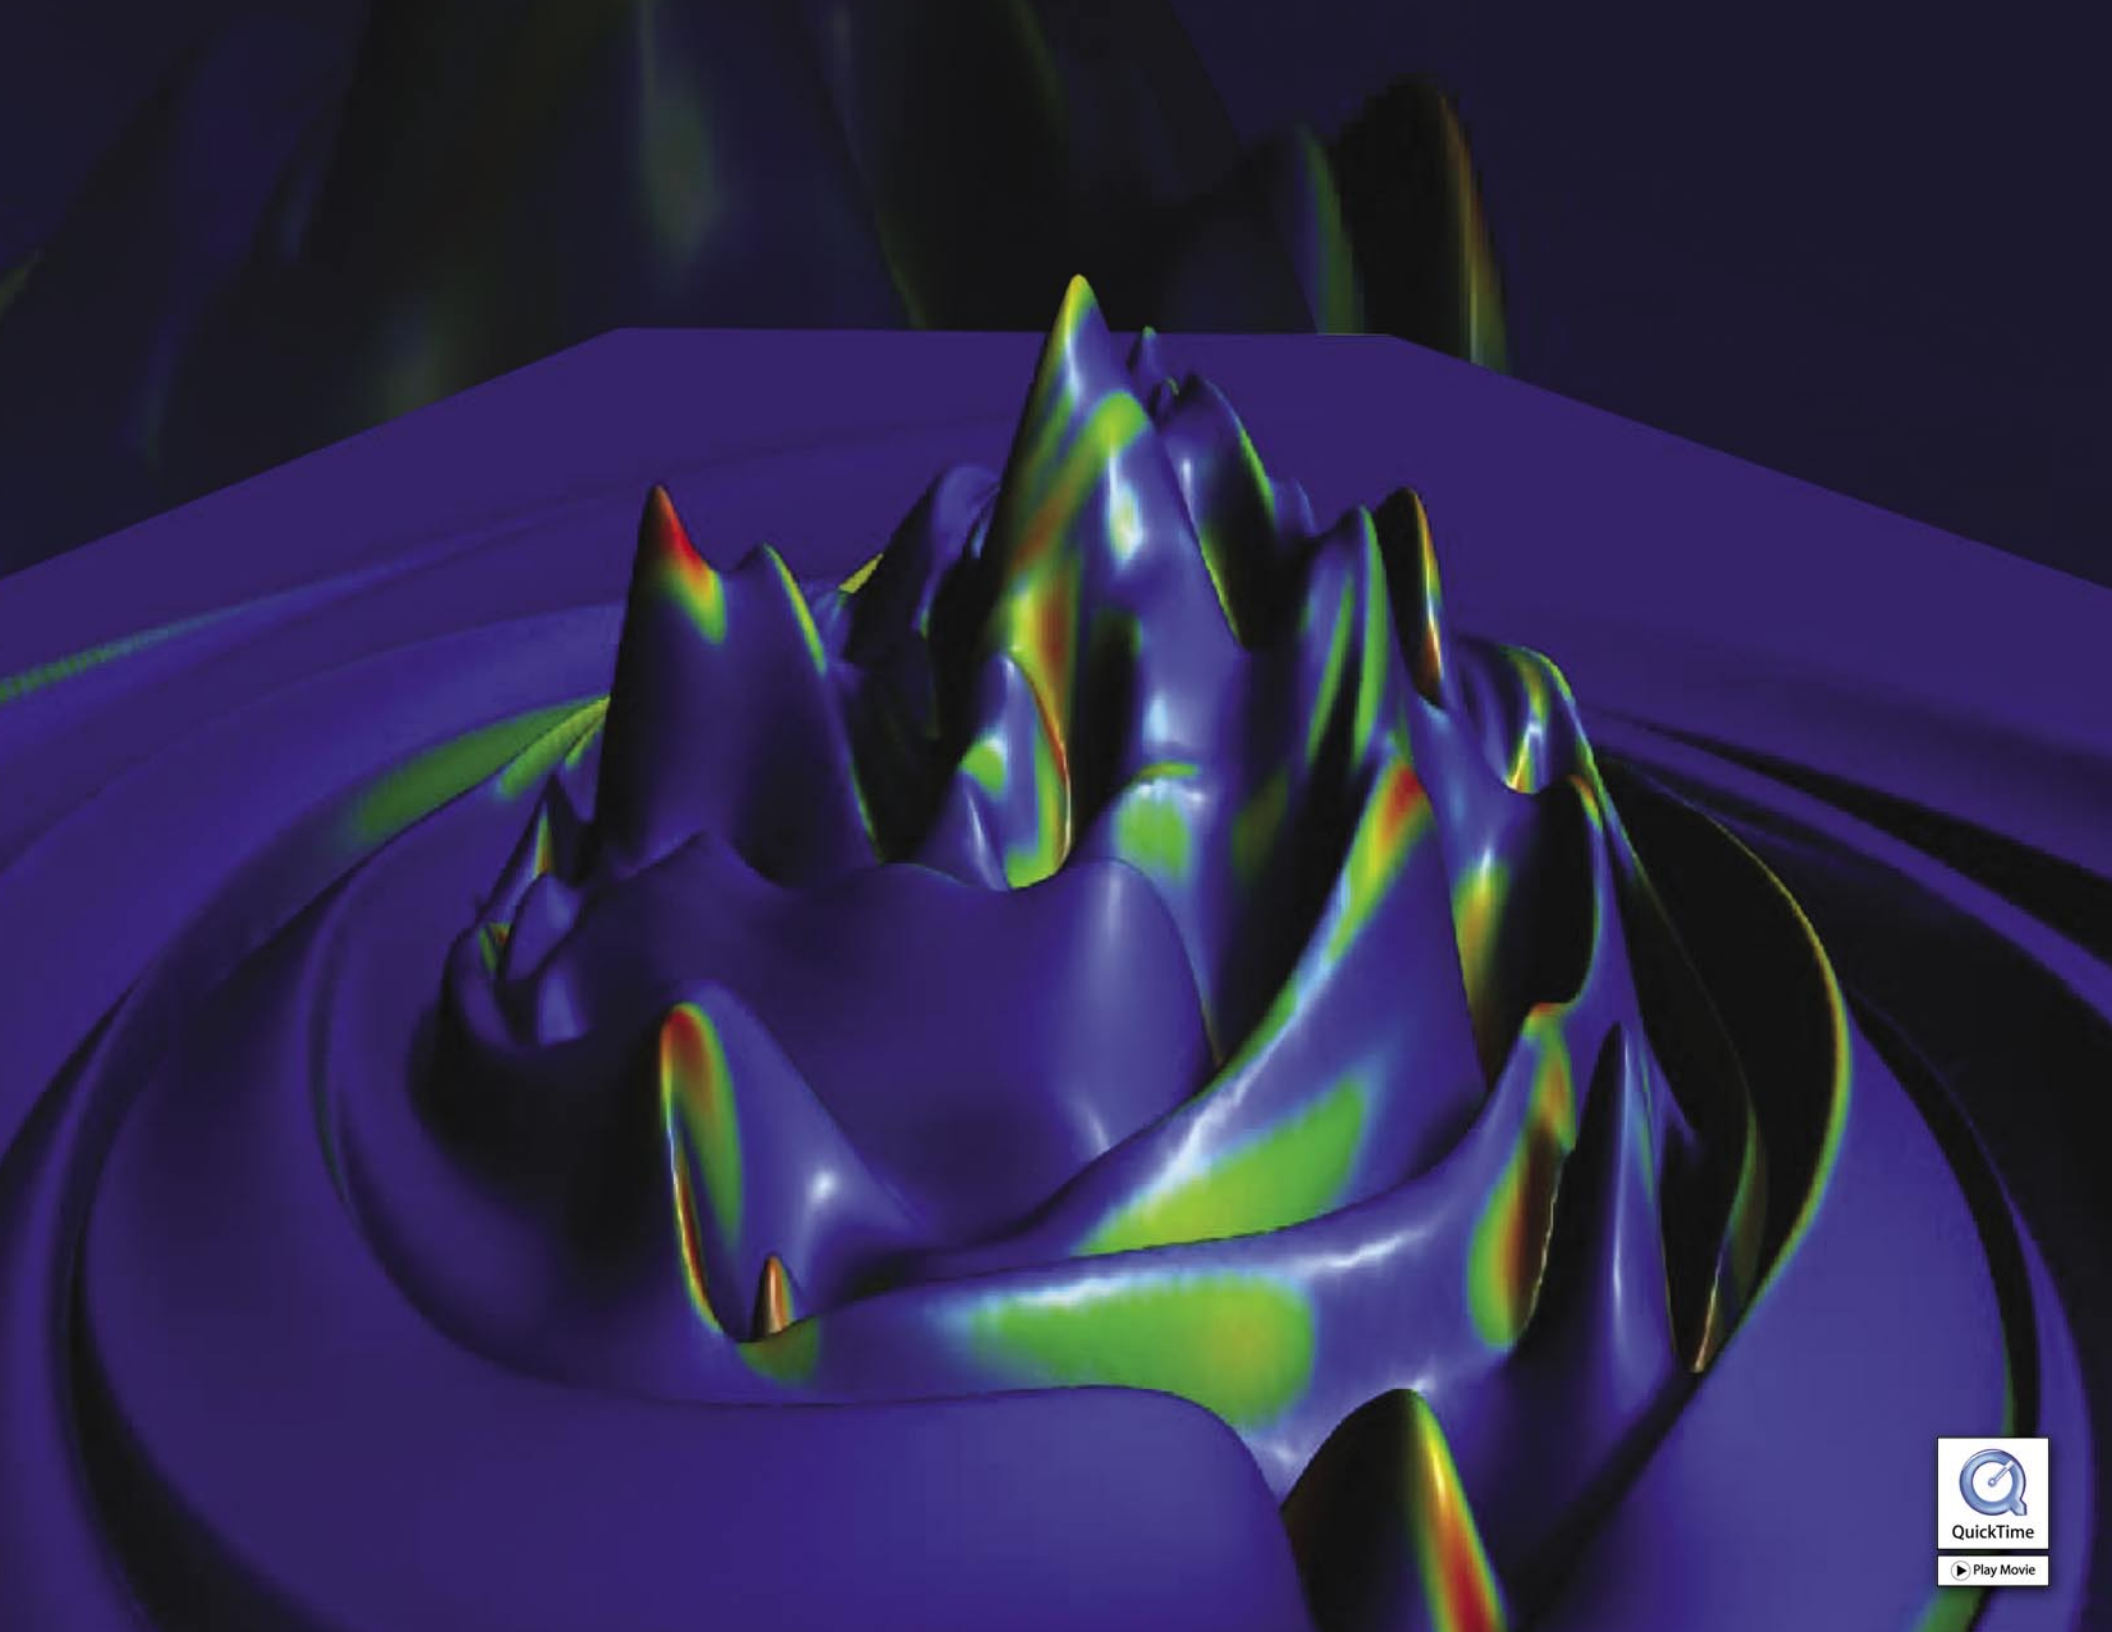
QuickTime
Play Movie

The Computation Directorate partners with every major Laboratory program in the development of software applications needed to meet programmatic milestones. One of the ways we do this is through the matrixing of highly-skilled computer science professionals who work side-by-side with domain experts to design, implement, and use these codes. The applications span a range of scientific and technical domains: simulation of fluid dynamics under extreme conditions; efficient, scalable, parallel multi-physics frameworks; computational modeling of materials; real-time control systems for large science experiments; and analysis of large sets of information from disparate sources.

Our Directorate employs more than 400 computer scientists and applications programmers, which makes us one of the largest applied computer science organizations in the world. We distinguish ourselves by combining outstanding discipline expertise with acquired domain knowledge, that is, an understanding of the application areas in which we work. This allows us to partner with the programs we support and helps to ensure that we provide innovative and relevant computing solutions to our customers. We strengthen these ties through the programmatic alignment of our application development efforts. This also has improved our responsiveness as a Directorate to Laboratory needs.

Figure 3.00-1 (facing page). The VIEWS visualization project simulated the dynamics of vortical flow, resulting in this height map of the vorticity (local fluid rotation) field.

The nature of the partnership between Computation and a given LLNL Program depends on the needs of that program, but the overall approach has many common features. Many programs rely on a small number of computer codes that grow and evolve over time (often decades) to meet specific needs. Computation may assign 2–20 computer scientists to each code to work with domain experts to define, implement, optimize and use these codes. The computer science expertise that must be brought to bear on each problem can differ dramatically from one problem to the next. For example, a stockpile stewardship application generally demands parallel computing and scalable algorithms expertise. These complex codes generate terabytes of data, the analysis of which requires sophisticated scientific visualization capabilities. In contrast, the real-time control system for NIF oversees every aspect of this unique experimental resource. This system has more than 60,000 control points and must manage timescales of events ranging from shot sequences (over hours) to laser pulse timing (trillionths of a second). Finally, in many applications, including bio-defense, the emphasis is on novel and domain-specific data analysis capabilities, for instance, the ability to identify unique DNA signatures via scalable pattern recognition techniques.

Underpinning all of our code development is the expectation that we will deliver high-quality, innovative, and maintainable software applications for our customers. We must develop a cost-effective approach that can be adapted to diverse programming situations, ranging from prototype research systems to mission-critical systems that cannot fail under any circumstances. Toward this end, we have spearheaded the Laboratory's new Software Quality Assurance initiative, to develop and promulgate risk-based, customizable software engineering practices. Finally, we strive to provide innovative software solutions by leveraging the Directorate's research activities (see Section 4). For example, our scalable algorithms research results are being integrated into ASCI and Office of Science applications.

The remainder of this Section highlights some of the work being done in Computation in partnership with LLNL Programs. Some of these projects are in the early stages of development, while others are reasonably mature and producing results. In each case, we try to highlight how we have applied our skills and expertise to contribute to solving the problems. We take pride in both our discipline capabilities and in the programmatic successes we enable through the timely and cost-effective use of these capabilities.

3.01 ASCI Code Development

Problem Description

The Defense and Nuclear Technologies Program relies heavily on large-scale scientific simulations in support of Scientific Stockpile Stewardship. These codes use a wide variety of physics and chemistry modeling techniques to predict and explain critical stockpile questions. We are an integral part of a team of physicists and engineers that solves problems of national importance.

Technical Approach/Status

Computation personnel assigned to each team contribute by developing, testing, optimizing, and enhancing these codes for a range of high-performance computing platforms. Overall, ASCI Code Development has had many advances this year in parallelization, algorithm development, and software quality assurance/software engineering (SQA/SE). This section will highlight two specific advances in parallelization and SQA/SE.

Progress in 2003

We have completed the initial parallelization of a Monte Carlo transport code that is showing excellent initial scaling results. The computer science advances include two options in parallelism that can work in combination. Spatial domain decomposition of the mesh and distribution of the particle load across the spatial domain allows for tracking a large number of particles in the code. Figure 3.01-1 depicts the two parallel modes. The verification on this code has shown it to compare favorably to other Monte Carlo particle transport codes on a subset of international criticality safety benchmark problems. The code is more flexible than similar Monte Carlo codes due to support for a variety of parallel run models, as well as capability to track several types of particles on problem geometries in 1, 2, and 3 dimensions. Finally, we have run hundreds of millions of particles on thousands of processors in 3D.

Several projects have developed and deployed enhanced automated testing and quality assurance software tools to handle multi-physics simulation codes and GUI testing. One testing tool, STAT quality assurance and performance monitoring system, continuously collects over 500 individual statistics from nightly regression tests and has found errors ranging from inefficient coding, to subtle errors in software and hardware platform changes. Tapestry and the Application Testing System (ATS) are regression-testing tools for parallel physics applications. Tapestry provides for MPI and OpenMP, batch and interactive testing, and several result comparison methods and reporting mechanisms. ATS provides a decentralized testing framework which allows someone who is an expert in a particular area to run a specific test in one or more ways while allowing someone who has no idea what the test does to run it without being told how. QTestViewer allows a developer working on multiple platforms to record GUI events, replay these events, save an event log, and then test the QT application systematically.

Significance

Using these enhancements and others developed by our team, we are now solving problems with higher quality and fidelity, more efficiently using valuable parallel resources, and increasing functionality and decreasing turn-around time for pre- and post-processing needs.

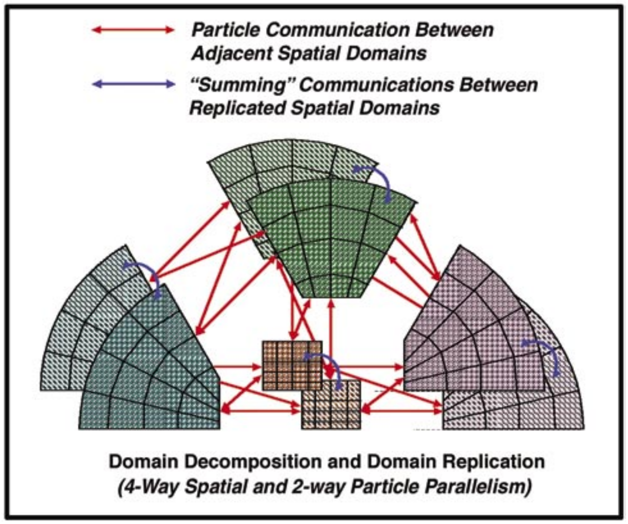

Figure 3.01-1. Domain decomposition and domain replication (4-way spatial and 2-way particle parallelism).

Problem Description

The mission of A/X Division's Verification and Validation (V&V) program is to provide scientific evaluation of the division's ASCI codes. A key aspect of the work consists of the design and implementation of analytic and semi-analytic tests that cover a broad range of physics. These tests are used to verify the physics algorithms implemented in the codes.

The results from at least one ASCI code are compared with results from legacy code, and with data from the National Ignition Facility, or other experiments that are selected on the basis of excellent data and are representative of key physics modules. One significant area of computer science support for this program is the development of a dynamic web site to report findings to the physics community.

Technical Approach/Status

The V&V web site is an interactive digital library that provides a means to publish internally the results from V&V studies. The web site is a modern web application utilizing Enterprise Java, commercial software, frameworks from the Apache Software Foundation, and emerging standards. At its core are an Apache web server running secure HTTP, an Oracle application server, and an Oracle database. The database serves as a repository for key parameters associated with a code run. Three categories of information are distinguished including biographical information about the run, key input parameters and figures of merit.

Progress in 2003

The physics reports, typically written in LaTeX, are converted to HTML and image files for browsing by using the tool Latex2Html. These are the only HTML pages in the system; the remainder of the content is generated by Java Server Pages (JSP) and servelets, which run as threads of execution under the control of the application server. The web site also incorporates a data browser integrated with the database. This is accomplished with an IDL session for on-line analysis of code results, and provides a graphical interface so that the user can visualize multiple data sets from database selections. In 2003, the web site became operational, with 10 reports and 350 runs registered in the database.

Significance

One of the key aspects of the web site is that a user can access results from a desktop by using any Web browser, as with Internet Explorer, for example. This eliminates a paper system and provides a means for rapidly distributing new information. The dynamic nature of the web site gives the user capabilities never before realized: the ability to interact with the report's data for both visual and analytic comparisons.

Figure 3.02-1. The W76 is one of the stockpile weapons for which capability has been demonstrated through verification and validation evaluations.

3.03 NIF Integrated Computer Control System

Problem Description

The National Ignition Facility, currently under construction at LLNL, is a stadium-sized facility containing a 192-beam, 1.8-megajoule, 500-terawatt, ultraviolet laser system together with a 10-meter diameter target chamber with room for nearly 100 experimental diagnostics. NIF is operated by the Integrated Computer Control System (ICCS), which will control more than 60,000 control points from a main control room (Figure 3.03-1). The control system is required to keep optical components precisely aligned over 1000 feet, and to finely orchestrate an automated shot sequence that takes place over hours, culminating in a billionths-of-a-second laser pulse that is kept in lock-step to a few trillionths of a second. The ICCS is constructed using a distributed object-oriented software framework that uses CORBA to communicate between languages and processors. This framework provides central services and patterns for building a layered architecture of supervisors and front-end processors.

Technical Approach/Status

The strategy used to develop ICCS calls for incremental cycles of construction and formal testing to deliver an estimated total of one million lines of code. Each incremental release allocates two to six months to implement targeted functionality consistent with overall project priorities. Releases culminate with successful formal off-line tests conducted by an independent Controls Verification and Validation (V&V) team in the ICCS Integration and Test Facility (ITF) and hardware integration labs. Tests are repeated on-line to confirm integrated operation and provide operator training in NIF. Offline tests in the ITF and in hardware integration labs, and these online tests in the NIF together identify 90% of software defects before the software is delivered to Operations. Test incidents are recorded and tracked from development to successful deployment by the verification team, with hardware and software changes approved by the appropriate change control board. Test metrics are generated by the verification team and monitored by the software quality assurance manager.

Progress in 20003

In 2003, NIF began its laser-commissioning program and has successfully operated the first four beams using the ICCS and by mid-2003 NIF had produced the highest energy 1ω, 2ω, and 3ω single laser energies in the world. All subsystems on NIF have been successfully fired for over 200 full system shots, achieving all scheduled project milestones. Approximately three-fourths of the NIF control systems software has been completed (including 250,000 lines of code in 2003) and used to commission and operate the first four beams of NIF.

Significance

Over the next several years, control system hardware commissioned on the first four beams will be replicated and installed to activate additional laser beamlines. Completing the remaining software is a large effort that involves completing shot automation software first for a bundle of eight beams and then for the remaining laser beams. During 2004, a separate testing effort is determining the performance limits of the control system and assuring the reliability needed to scale the control system to operate multiple bundles, and eventually 192 beams. The ICCS team is also structuring the higher-level server and shot automation software to readily meet the performance requirements as the laser is built out. This straightforward scaling flexibility is extremely important for the successful and reliable operation of NIF and was a key design goal when CORBA was chosen as the distribution mechanism for ICCS.

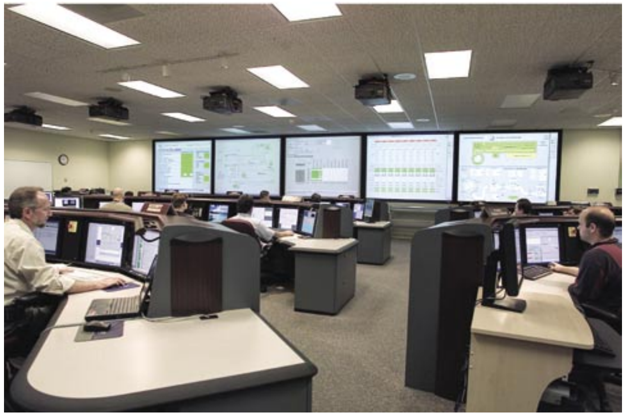

Figure 3.03-1. NIF is controlled from a main control room using the Integrated Computer Control System.

Problem description

Rapid identification of natural and genetically engineered pathogens is vital to our national security. Previously, efforts had concentrated on developing assays to detect organisms by targeting small portions of their genomes (e.g., single, well studied genes). The need arose for new tools that could rapidly identify multiple unique and specific regions of the genome, increase specificity, and reduce cost of development and time to deployment. This work is a continuation of a collaboration with BBRP biologists that counts the Human Genome Project among its successes.

Technical Approach/Status

Since August 2000, LLNL has pioneered the computational design of DNA signatures for detection of natural disease outbreaks or bioterrorism.
Our DNA signature pipeline is a fully automated software system for identifying unique regions on pathogen genomes and selecting optimal sequences for development of real-time detection assays (Figure 3.04-1). These signatures are stringently tested by LLNL biologists and by external collaborators. They are validated by the Centers for Disease Control (CDC) and placed into regular use for environmental monitoring. Air samples are collected on filters, and then analyzed for the presence of pathogens. Results can be obtained in as little as 90 minutes. As of early 2004, our assays have more than 700,000 live bio-defense uses with zero false-positives.

Progress in 2003

We scaled our original DNA signature pipeline to run in parallel on a 24-CPU server. We can now process bacterial pathogens (e.g., Anthrax) in less than two hours. An invitation from the CDC to provide signature design help led to signatures for smallpox, monkeypox, and SARS—all in response to unplanned world events. Our system now automatically downloads new and updated microbial sequences from multiple public Web sites and checks all fielded signatures for potential false positive/negative results. We linked the DNA signature pipeline to a fully automated annotation system, which provides us with the additional capability of down-selecting from among many good signature candidates based on their association with genes of interest (e.g., virulence factors).

Leveraging our success with DNA signatures, in 2003 we prototyped a protein signature pipeline that can locate unique protein-based diagnostic targets on a whole-proteome scale. Our system identifies unique peptides, locates them on 3D protein models, and annotates for best-choice selection.

Significance

In 2003, our signatures entered BioWatch, which daily monitors the air in several dozen US cities. We have been asked to apply our computational system to a variety of human, animal, and plant pathogens. Success of our DNA signature design effort has established LLNL at the forefront of pathogen diagnostics. Analyses we performed in 2003 for the CDC have directly influenced the nation's policy on pathogen defense.

Figure 3.04-1. Hand-held detectors are important as the first line in Bio-Defense and pathogen identification.

3.05 Radiation Transport for Cargo Inspection

Problem Description

The Radiation Transport project develops a deterministic neutral particle transport simulation tool to model the time-dependent and steady state transport of neutrons and gamma rays through materials. This tool is being provided to DHS researchers to help in the design of active and passive radiation detectors used for detecting special nuclear material in cargo at inspection sites. It will provide them with a deterministic modeling capability that complements current Monte Carlo (MC) tools, is potentially much faster, can calculate solution sensitivities, and is in-house. Such a tool could be used in portable radiation detector configurations for DHS.

Technical Approach/Status

Neutron and gamma ray transport will be simulated in 1D, 2D, and 3D geometries, and will include delayed neutrons and gamma rays produced by fission, when appropriate. In many problems of interest to DHS, the current MC methodology is computationally intensive, requiring long run times. In such instances, an equivalent 1D or 2D deterministic capability would require much less work, and be extremely useful to DHS researchers. This is particularly true for simulations in highly diffusive media. Simulations in 3D may also be important. An added benefit is the ease with which solution and/or detector sensitivities with respect to design parameters can be calculated when using deterministic methods.

Progress in 2003

During 2003, we added the capability of solving the neutron kinetics equations to the transport code Ardra. Figure 3.05-1 shows a cross-section of a simple model of a cargo container containing a highly enriched uranium target (purple) shielded by simulated cargo (plywood, in green), a localized source (on the right), and a detector (black). Figure 3.05-2 is a snapshot of a time-dependent simulation showing a neutron pulse in the scalar flux of the 14MeV neutron energy group at 46μs as the pulse is traveling through the target.

Significance

Modeling is a significant part of the design process for active and passive radiation detectors. Providing a tool such as the one described above will allow for faster and more effective detector designs. The ability to track the delayed neutrons that result from fission is crucial to accurately model the delayed detector response seen in an actual experiment.

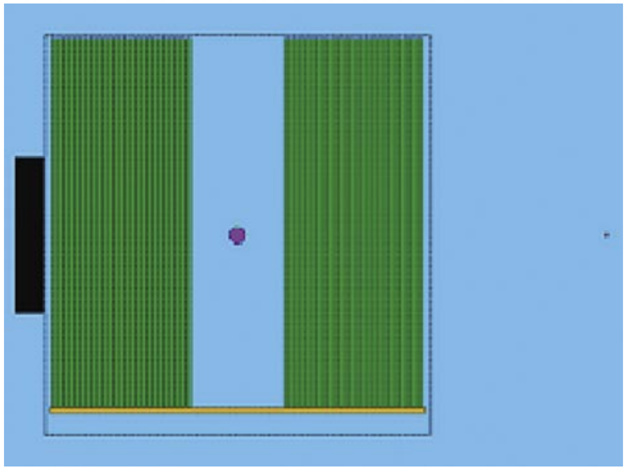

Figure 3.05- 1. A typical cargo container with a hidden target containing high-Z material.

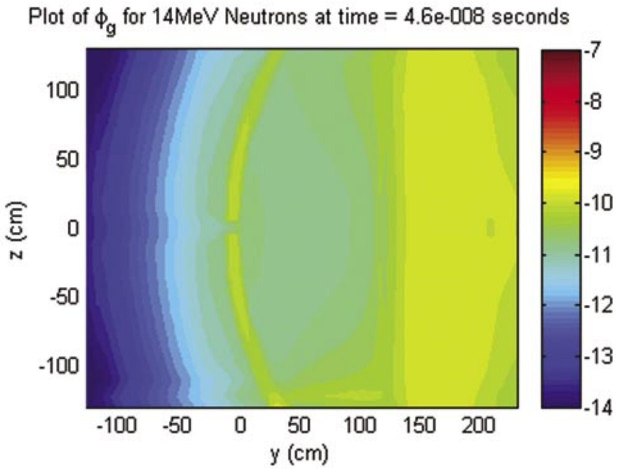

Figure 3.05-2. Scalar flux can be mapped as a pulse travels through the target.

Problem Description

The National Atmospheric Release Advisory Center (NARAC) provides to emergency managers the tools and services that map the probable spread of hazardous material (nuclear, radiological, chemical, biological, or natural emissions) accidentally or intentionally released into the atmosphere. Customer growth over the past few years has been significant—NARAC now supports many federal agencies, including DOE and DHS, as well as a growing number of state and local government agencies. To facilitate this increased external customer access, new systems and tools were developed, including the NARAC Enterprise System (NES), the NARAC iClient, and the NARAC Web. The NARAC software development team is composed primarily of Computation personnel.

Technical Approach/Status

NARAC is a distributed system, providing modeling and geographical information tools that run on an end-user's computer system, as well as real-time access to global meteorological and geographical databases and advanced three-dimensional model predictions from the NARAC Central System (NCS) at LLNL. The NCS is an object-oriented system, written in C++ and Java, using CORBA for communications and an OODBMS for object persistence. The iClient is written in Java and communicates with NES using SOAP. The NARAC Web is dynamic HTML that uses HTTPS to communicate with NES. The NES is written in Java, is built on J2EE, and uses a JDBC-compliant database; NES communicates with NCS using CORBA.

Progress in 2003

The NES and Web were designed, developed, and deployed within one year. Security, flexibility, ease of use, and future expansibility were the main design goals. User authentication, encrypted data communications, and fine-grained security were essential. For example, fine-grained security ensures that each user has access only to the capabilities and information for which he or she is authorized. By the end of 2003, there were more than 500 NARAC Web users at all levels of U.S. government.

Significance

The full NARAC System was used to support major exercises, alerts, and potential emergencies. For example, the Web and iClient were used extensively during TOPOFF2, the largest national emergency preparedness exercise since the terrorist attacks of September 11th. Set in Seattle, this exercise simulated a "dirty bomb" and involved emergency personnel from the city, county, and state governments, as well as 19 federal agencies, including DOE and DHS. Virtually all participants used NARAC predictions during this exercise. The Seattle Hazmat team and Incident Commander on scene used wireless communication and laptop-based NARAC iClient to submit and access NARAC predictions. Plume predictions were distributed using the NARAC Web to Seattle Fire and EOC and other county, state and federal agencies in real-time. Officials from the Mayor of Seattle, to the DHS Secretary, to White House personnel were briefed using NARAC predictions distributed over the Web.

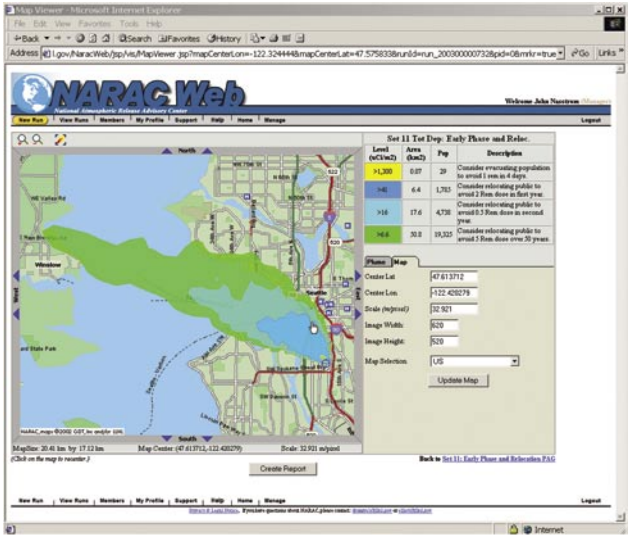

Figure 3.06-1. Sample NARAC Web plot from the TOPOFF2 emergency preparedness exercise.

3.07 Coupled Climate and Carbon Modeling

Problem Description

Contemporary climate models use prescribed carbon dioxide (CO_2) concentrations to predict the resulting climate. To properly assess the impacts of fossil fuel burning, however, one must instead base the climate computation on anthropogenic (human-induced) emissions of CO_2. Computation and Atmospheric Science Division (ASD) personnel have collaborated on the development of an integrated climate and carbon-cycle model that predicts the fate and climatic effects of fossil fuel-derived CO_2 and are applying it to analyze global warming and other related effects through the 21st century and beyond. This is the most comprehensive, and first American, fully coupled climate–carbon simulation system.

Technical Approach/Status

In this collaboration, Computation personnel have lead responsibility for the enabling technology, and ASD personnel for the scientific study. Computation members enhance and couple together the relevant component codes, parallelizing where necessary to create a scalable model that can execute on a multitude of high-performance architectures. In particular, they re-partition the land points in the terrestrial biosphere model to balance the computational load. This requires the institution of high-speed transposes to connect the terrestrial biosphere and atmospheric model domain decompositions. Researchers take advantage of both distributed and shared memory parallelism where possible.

Progress in 2003

With the model integration largely completed previously, 2003 was devoted to scientific simulation. Two main studies were completed, each involving several multi-century simulations, and each simulation taking roughly 50 days of active wall clock time. In the first study, project members analyzed the effects of CO_2 fertilization on the atmospheric concentration of CO_2. Depending on the extent to which CO_2 fertilization saturated with increasing CO_2 levels, this ranged from a doubling to a tripling of atmospheric CO_2 levels over the course of the 21st century (Figure 3.07-1).

In the second study, researchers varied the sensitivity of the climate to the radiative forcing of CO_2 and saw surface temperature increases of 3° to 8°K (Figure 3.07-2). While the project team believes these runs bracketed the degree of anthropogenic global warming, the variations underscored the importance of accurate, comprehensive modeling.

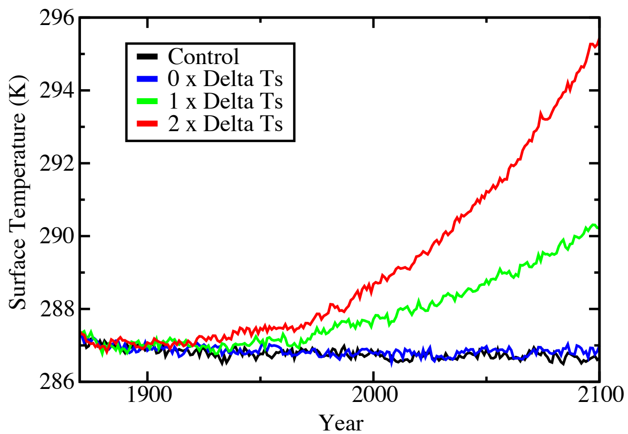

Figure 3.07-1. Simulated atmospheric CO_2 from 1870 through 2100, as a function of CO_2 fertilization (control run in black, standard case in green, saturated case in red). Saturated fertilization gives rise to a larger increase in CO_2.

Significance

The experience and expertise of Computation personnel in high-performance computation, parallel code design, and modification and integration of large scientific programs has been instrumental in the success of this highly relevant endeavor. Having the first integrated climate and carbon modeling capability is an important step toward fulfilling DOE's mission to reduce uncertainties arising from climate–carbon feedbacks so that we can better address scientific and policy-related questions involving the climatic effects of burning fossil fuels.

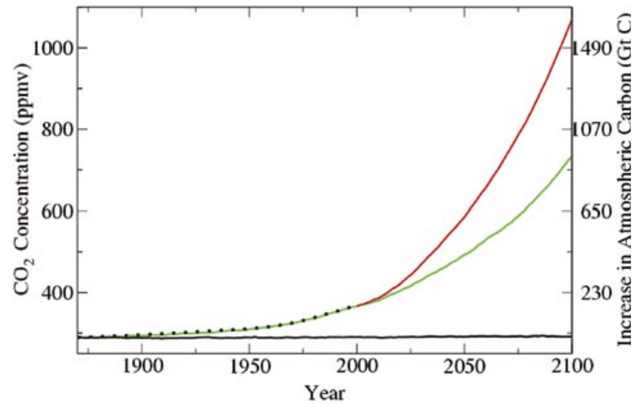

Figure 3.07-2. Evolution of globally averaged surface temperature, as a function of climate sensitivity to radiative forcing (control run in black, standard case in green, highly sensitive case in red). The expected increase in average surface temperature is anywhere between 3° and 8° K.

Problem Description

The properties of condensed matter at high pressure are both difficult to measure, and difficult to derive from what is known at atmospheric pressure. In particular, microscopic structural properties, (e.g., the type of chemical bonds present) and physical properties (e.g., the melting temperature) can be drastically modified when a substance is subjected to pressures reaching millions of atmospheres (Megabars). An experimental determination of these properties is often complex and expensive.

Technical Approach/Status

Recent progress in the technology of First-Principles simulations provides a new avenue for the exploration of properties of condensed matter in extreme conditions. First-Principles simulations are based on fundamental properties of matter derived from quantum mechanics, and do not rely on any empirical or adjusted parameters, thus providing a genuine theoretical prediction tool. The GP First-Principles simulation code developed in Computation has been used over the past years at LLNL to study the properties of fluids at high pressure.

Progress in 2003

In 2003, an important new capability was added to this simulation method by combining it with the so-called "two-phase" simulation approach. For the first time, it was possible to simulate accurately the solid–liquid interface of a molecular substance at high pressure and high temperature with a First-Principles approach. Using this new method, LLNL scientists in the Computation Directorate and in the Physics and Advanced Technologies (PAT) Directorate were able to predict the melting properties of lithium hydride up to a pressure of 200 GPa (2 Mbar). Results were published in *Physical Review Letters*.

Significance

This new type of simulation reaffirms and extends the role that First-Principles simulations will play in exploring the properties of condensed matter in extreme conditions. Being based on First Principles, this method is applicable to any other substance as well. The current parallel implementation, coupled with the powerful new high-end computing platforms available at the Laboratory, will further strengthen the position of LLNL at the forefront of high-pressure simulation research and Equation of State (EOS) calculations.

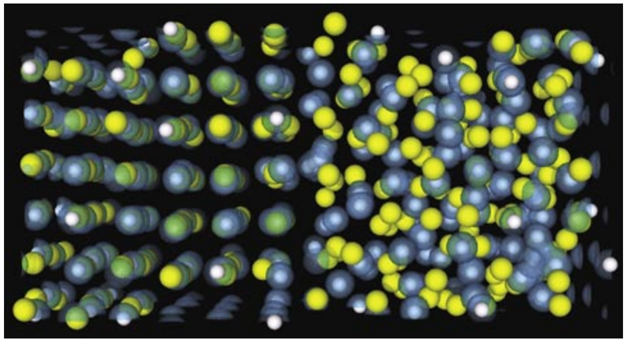

Figure 3.08-1. Two-phase First-Principles simulation of lithium hydride showing the coexistence of a solid phase (left) with a liquid phase (right). The simulation cell contains 432 atoms.

3.09 Computer Incident Advisory Capability

Problem Description

The Laboratory's Computer Incident Advisory Capability (CIAC) is a national DOE/NNSA cyber analysis center. CIAC notifies the Complex of vulnerabilities being exploited, recommends countermeasures, provides an overview of the current attack profile, and assists sites. CIAC focuses on specific threats and malicious activity targeting DOE/NNSA, and is developing a predictive analysis capability. This Advanced Warning and Response System (AWARE) assists DOE in preventing incidents rather than simply reacting to them after the fact.

Technical Approach/Status

AWARE strives to provide "Information-to-Insight-to-Action." It integrates technologies such as data mining, pattern recognition, statistical trending and traffic analysis, attacker profiling, and advanced visualization techniques. The system searches for trends and indications of possible attacks by analyzing potential threats, sites' sensor data, and data supplied by outside sources. AWARE will be able to pinpoint the right data at the right time, integrate sensor networks and databases into 24/7 operations, and allow actionable correlations to be made in near real-time.

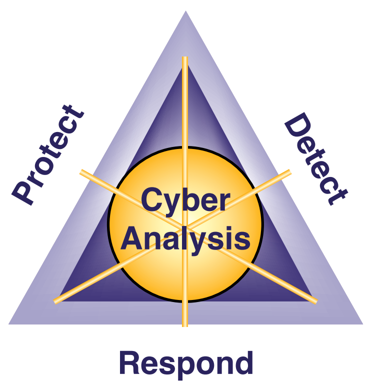

Progress in 2003

Network traffic is analyzed within 15 minutes of receipt, by streamlined processing of 2 GB/day of data collected by sensors deployed at several DOE/NNSA sites. Statistics are calculated in parallel threads, synchronized, and loaded into an Oracle table for reporting and visualization.

The AWARE portal (alpha release) disseminates results of the hourly analysis. The portal provides authenticated access to site-specific information and non-site–attributed Complex-wide information to DOE/NNSA security personnel.

Clustering techniques profile IP addresses for normal behavior, so that current activity can be compared and aberrant behavior detected. This reveals key variants that distinguish types of behaviors that leave cyber fingerprints.

Significance

As cyber attacks continue to rise in sophistication and virulence, cyber indications and warning systems are more critical than ever. Vulnerability exploitation time is decreasing dramatically, while the cost of repairing the damage has doubled each year from 2001 through 2003. The sooner new exploits or vulnerabilities are detected, the earlier DOE/NNSA can take action against them.

AWARE provides DOE/NNSA with an effective cyber indications and warning capability. It proactively protects Departmental assets from compromise, thus averting potential incidents and their ensuing impacts on productivity throughout the restoration and recovery periods. Through AWARE, CIAC advances the current "Protect-Detect-Respond" security defense strategy to one facilitated by anticipating adversary attacks, assessing intrusions, and assisting the sites in adapting their security architectures to proactively counter the attack. CIAC's work in this area allows DOE/NNSA to advance technology beyond intrusion detection to intrusion forecasting.

Figure 3.09-1. CIAC provides 24/7 incident response, intrusion analysis, vulnerability response, and counterintelligence support.

Problem Definition

The Information Operations and Assurance Center (IOAC) is a capability with a focus on information operations, knowledge management, and analysis tool development. IOAC's information management and analysis tool development is a noteworthy strength that we are aggressively applying to several major homeland security mission areas including infrastructure protection and bio-defense analysis—the current program is handling and correlating diverse information feeds from multiple sources and creating a graph that interrelates the information from these sources.

Technical Approach/Status

The Information Fusion and Analysis area has been making progress in building large-scale information fusion systems. We have developed semantic graph-based technology to perform real-time threat analysis and warning. The semantic graph facilitates this by extracting important relationships and correlations from a plethora of diverse data sources. The Network Analysis Tools area is dedicated to development of science and technology solutions in support of the analysis of networks. This focus area develops tools to automatically build a network model, graphically visualize it, and analyze it for attributes and patterns of interest including identification of vulnerabilities.

Progress in 2003

Enhanced network analysis tool capabilities fused disparate data with network-related information and added a GIS front end to the mapping system. We initiated a collaborative program to develop the Information Fusion and Analysis technology for the DHS. Additionally, we completed the design and partial development of a prototype system for DHS to make inferences from diverse data sources and data types.

Significance

The 2003 deliverables provided new capabilities for the program sponsors. The information fusion engine will enable the DHS to automate the process of "connecting the dots" across numerous diverse data sources and data types. The system security infrastructure supports U. S. law and privacy requirements.

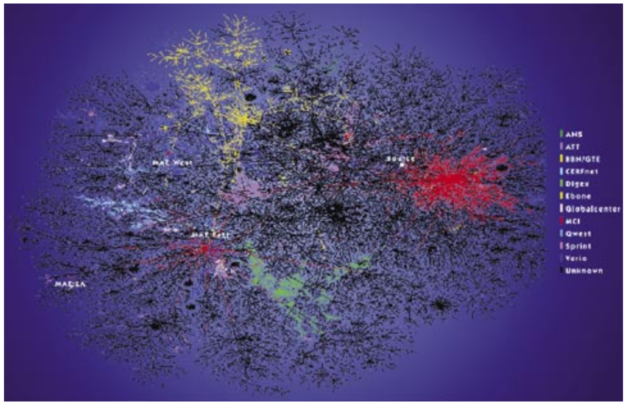

Figure 3.10-1. Network Analysis Tools automatically build, graphically visualize, and analyze a network model for attributes and patterns.

3.11 Software Quality Assurance

Problem Description

Outstanding software quality helps projects achieve their performance, cost, and schedule parameters. Each LLNL directorate has evolved its own Software Quality Assurance (SQA) approach, with varying degrees of documentation for those practices. In the absence of a consistently implemented Labwide SQA policy, we run the risk that external standards could be imposed on LLNL. Similarly, the Defense Nuclear Facility Safety Board has raised concerns that inconsistent practices might reduce confidence in the safety systems and safety analysis performed across the DOE Complex. These factors have prompted DOE and NNSA to raise the formality level of SQA for software used in critical applications.

Technical Approach/ Status

Led by Computation, software engineers from each directorate are collaborating to develop an institutional SQA Implementation Plan (SQA IP). LLNL will employ a multi-tiered, risk-based, tailorable approach to defining the SQA practices that apply to each project. Each project will grade its software components based upon the risk associated with their use. The graded approach will balance the hazards associated with the work the software is performing with the degree of rigor in the software quality practices. The SQA IP team will define various tiers of risk, and designate the software quality practices appropriate for each.

The IP will define these risk categories and associated practices based upon a set of over-riding SQA principles. The SQA IP is a blended model of external standards and current LLNL SQA best practices, to ensure added value to existing and future projects. In addition to clarifying appropriate SQA practices, this institutional approach satisfies the expectations of the external oversight organizations that increasingly scrutinize our software quality.

Progress in 2003

As the institutional approach is being developed, client programs have continued to enhance their own software engineering practices. ASCI sponsored a two-day, Tri-Lab workshop to share information about software engineering best practices, so that agencies might learn from one another. NIF continued to focus on its internal software engineering practices, and retained Raytheon to establish a unit-testing organization. NARAC developed an automated test suite to streamline its process. ARGUS increased its testing coverage with integration testing during the development phase, conducted by its own SQA team. This early testing protocol provided an earlier defect detection.

The Directorate also developed a series of seminars and training sessions to focus on software engineering practices. The seminar series for fiscal year 2004 includes usability, secure coding, testing, configuration management, risk management, and software quality updates. In an effort to better understand our customers' needs, Computation has begun to inventory current practices and tools used by Lab personnel in developing application software.

In 2003, six Computation employees became certified software quality engineers through the American Society for Quality. This was made possible by Computation's sponsorship of the software quality training and certification testing onsite. The Directorate is also hiring trained software quality experts to meet increased program demands. We are working closely with DOE in developing their SQA practices by stationing a LLNL employee at Headquarters.

Significance

Computation's SQA continues to enhance and enable improved software engineering practices Labwide. With Computation leading the effort, the current LLNL institutional SQA policy was approved on October 3, 2003. The primary policy objective was and is to deliver the best software products to our customers, and thereby to create confidence in the performance, cost, and delivery schedules. By developing a consistent software quality framework, we ensure that we will meet customer needs and comply with their software quality requirements.

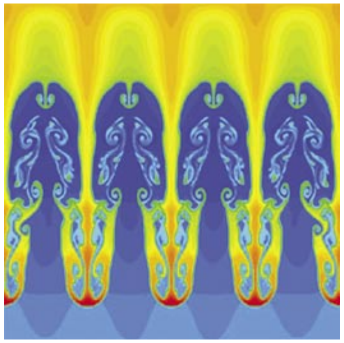

Figure 3.11-1. ASCI at LLNL relies on software quality assurance to produce accurate simulation results such as this hydrodynamic instability calculation.